Data Intensive Systems (DIS)

Data Intensive Systems

Dr. José L. Muñoz

DARPA/ITO

Outline

- The Data Intensive Problem & Opportunity
- Defense Data Intensive Requirements
- Goal
- Approach
- Product & Impact
- Roadmap

Today: Disk-Based Processing

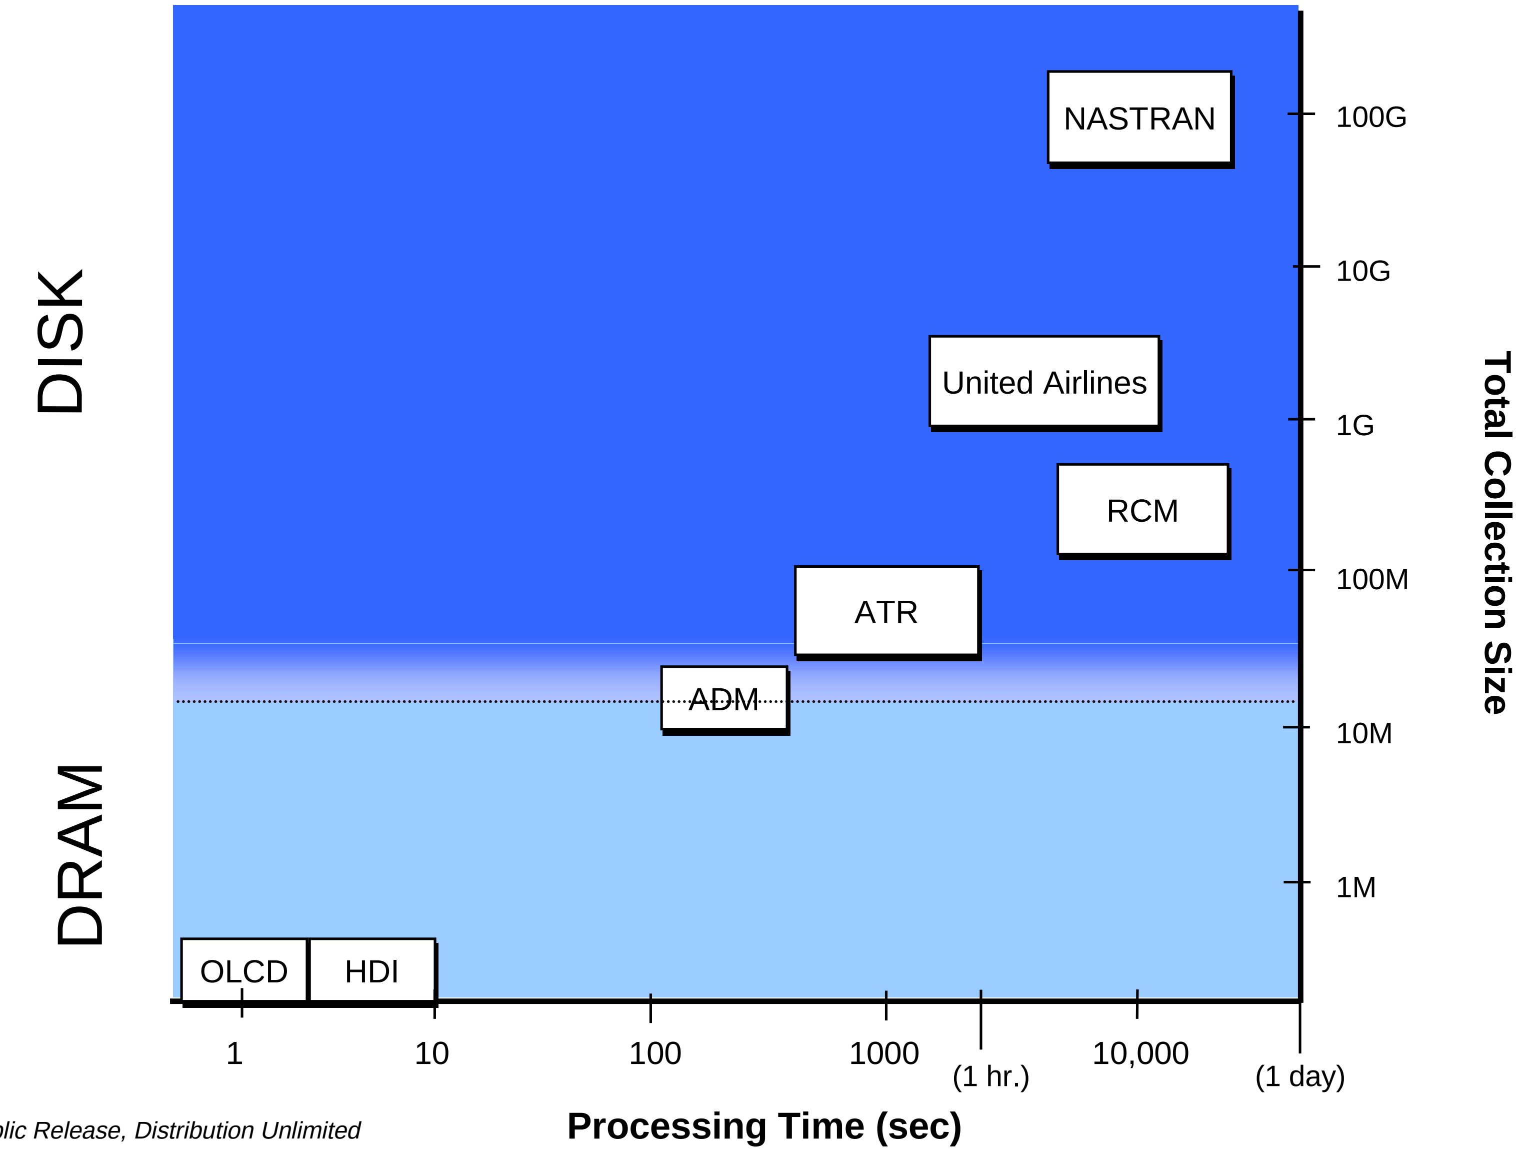

Opportunity:
Memory Based Processing
(Yr. 2005+)

Enable new regime of realtime data-intensive applications

Hours ➔ Seconds

Do for Memory systems what RISC did for processors

Realize total improvement of 3 orders of magnitude over existing systems

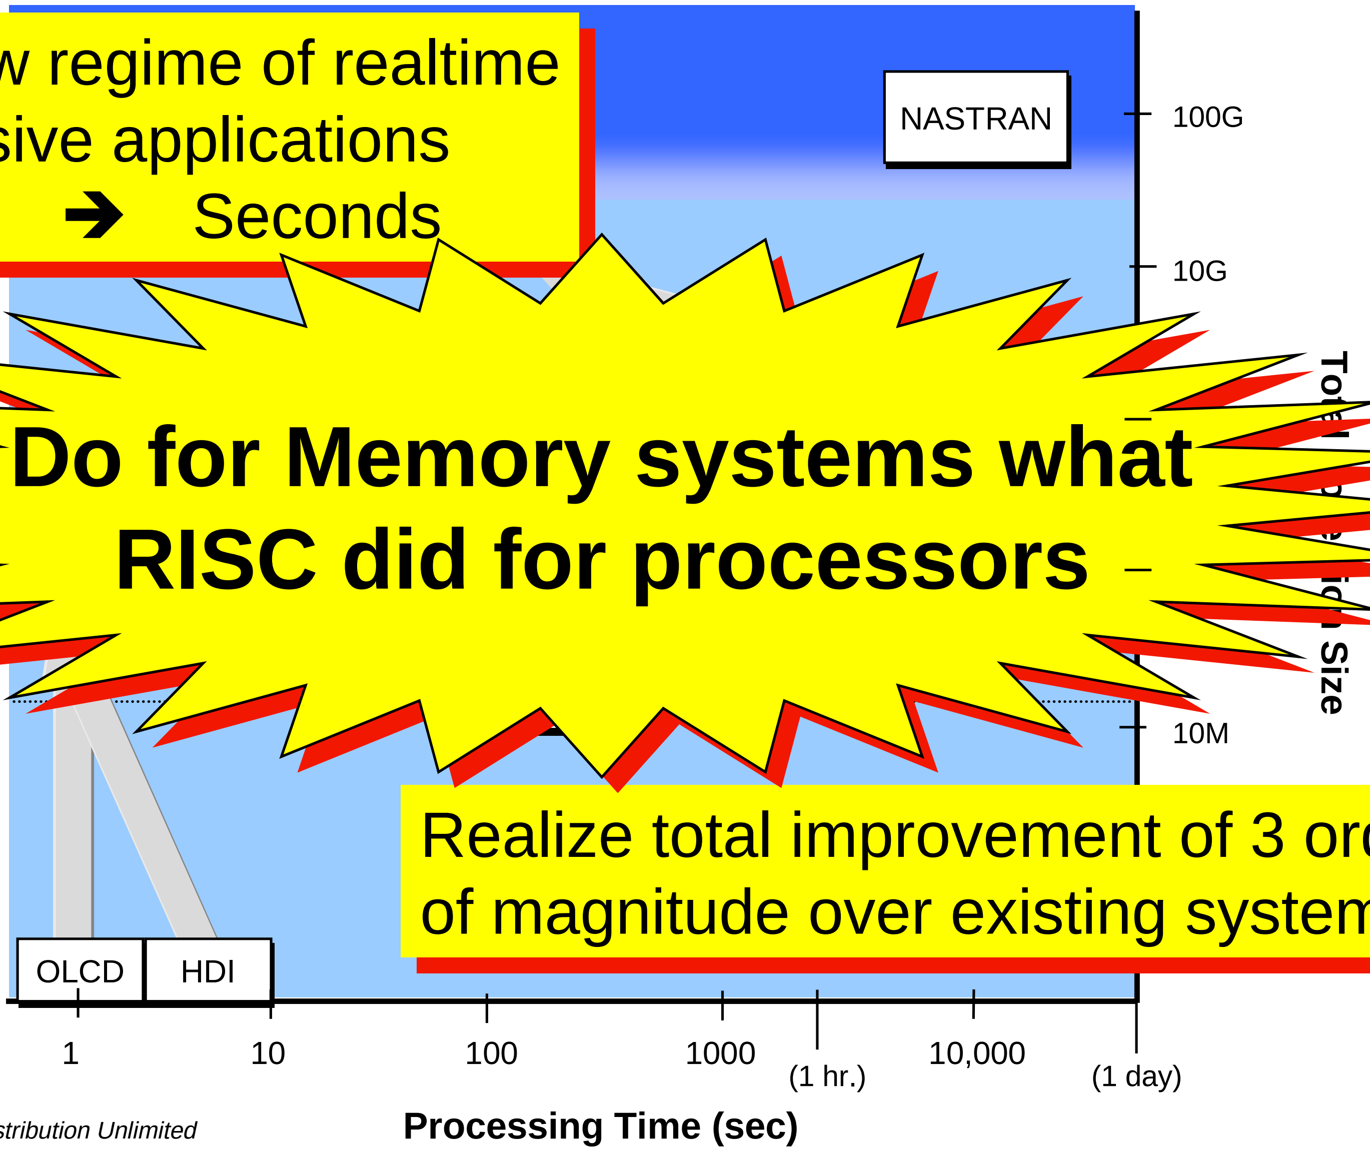

Problem: Logic and Memory Speeds Diverging

Memory Capacity

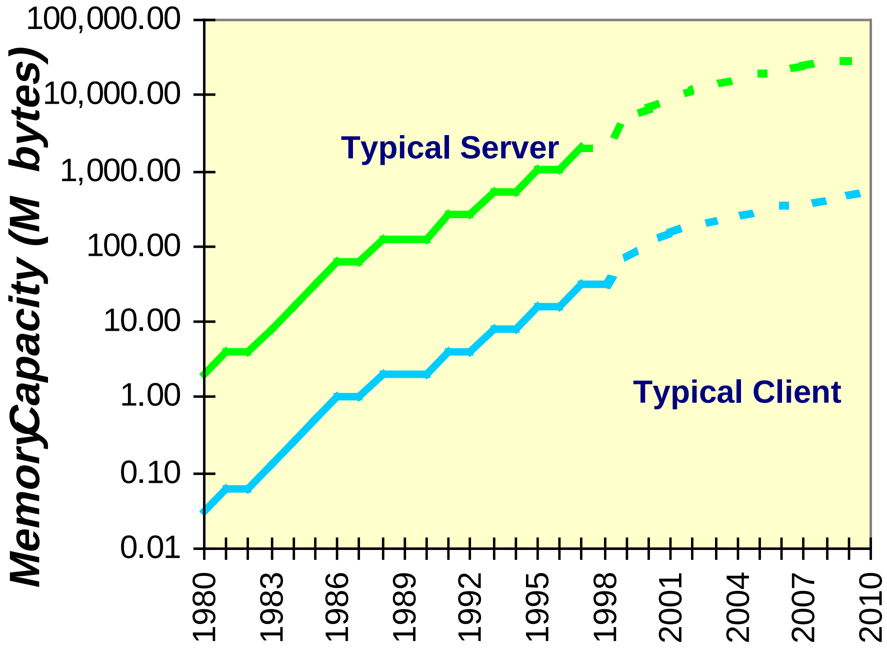

Memory vs. CPU Cycle Times

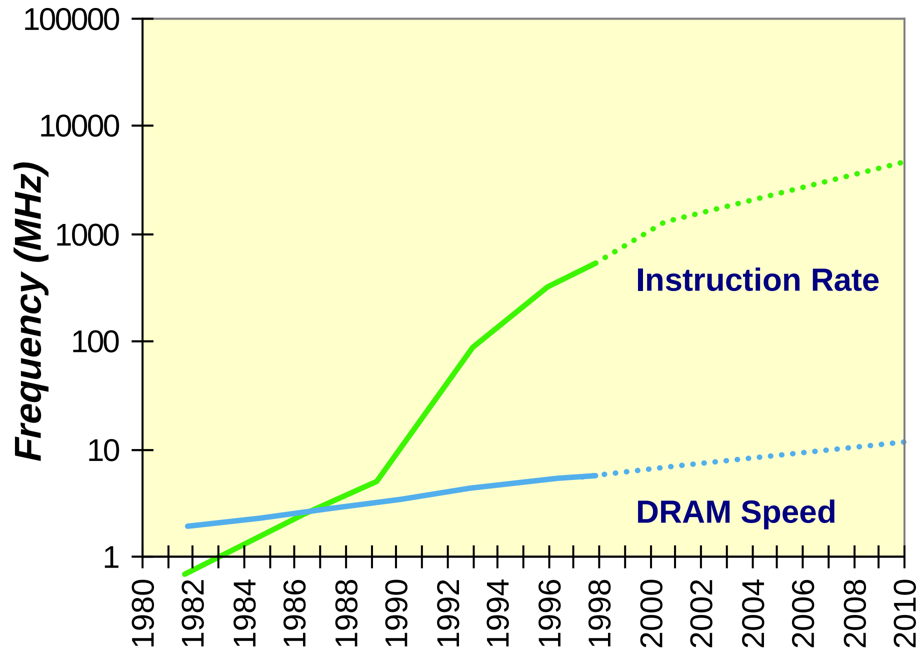

Memory access times has failed to keep pace with memory growth

Implications for In-Memory Databases

Memory Capacity

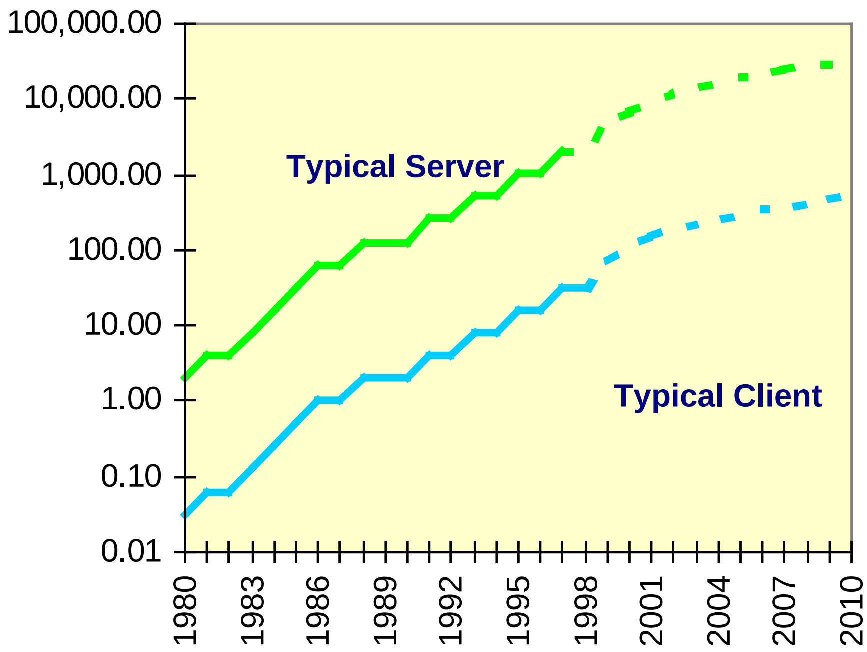

Query Capacity

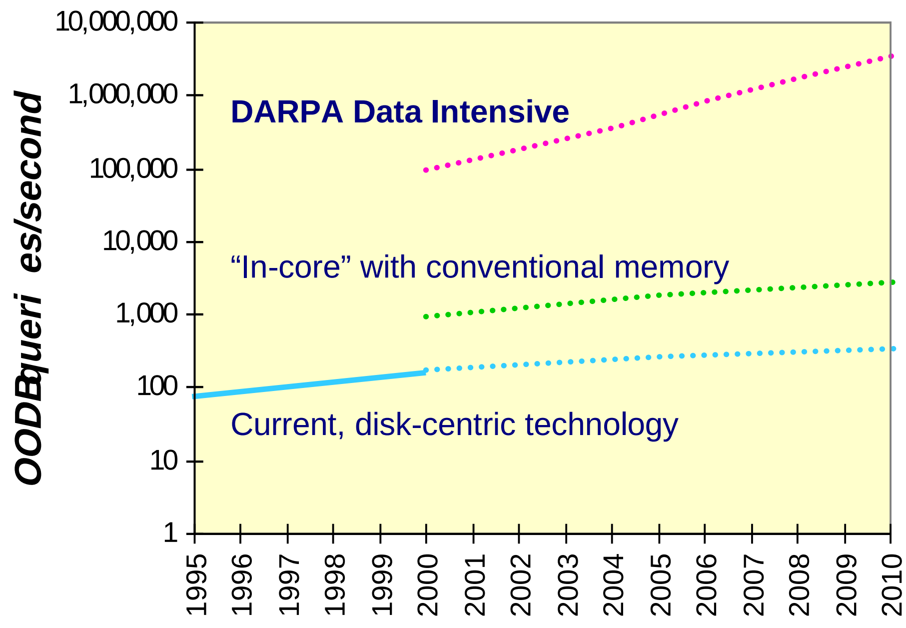

Programs/Data will grow to fill-up all of available memory

Implications for In-Memory Data Scan

Memory Capacity

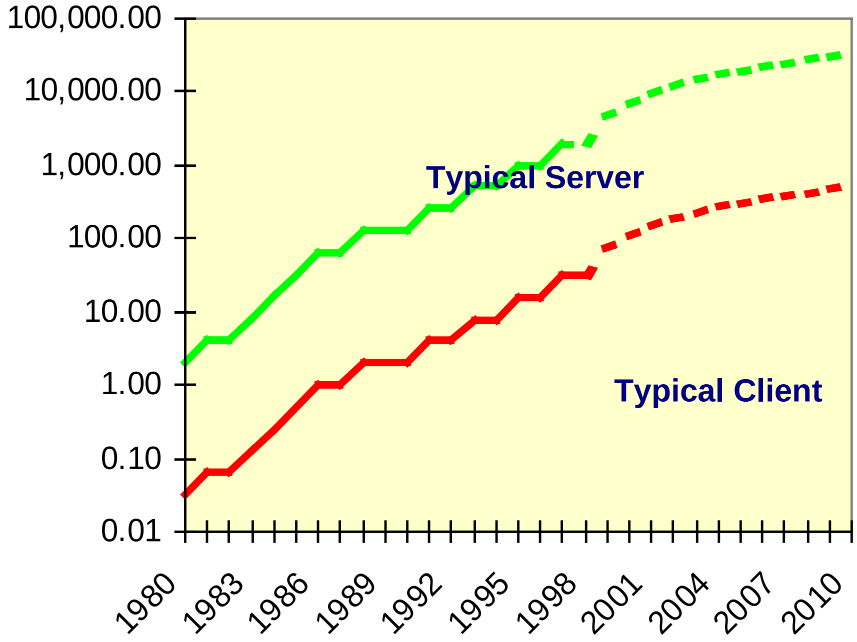

Scan Frequency

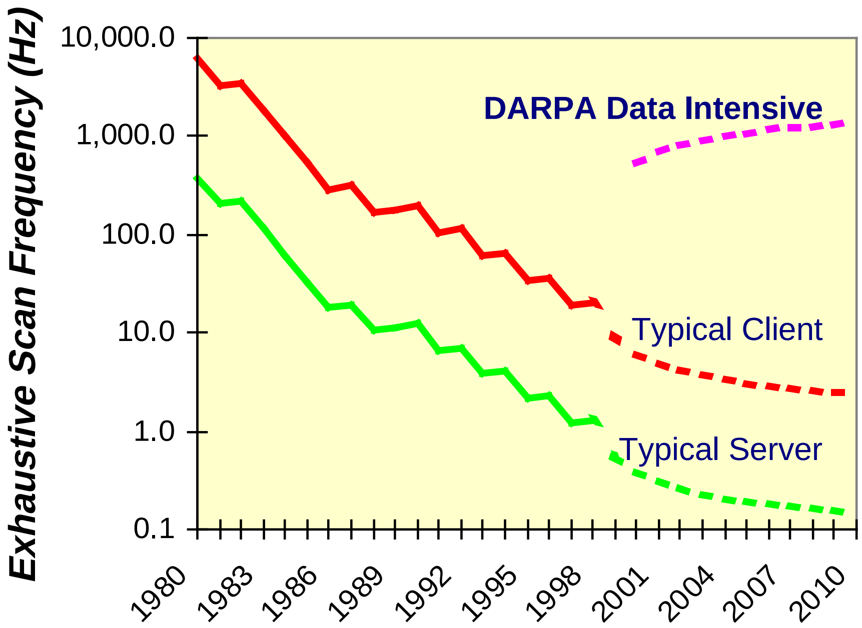

Metric: Frequency at which an object collection can be exhaustively scanned

Aggravating the Problem …

Traditional Premise:

Applications Have Small Working Sets of Contiguous Data

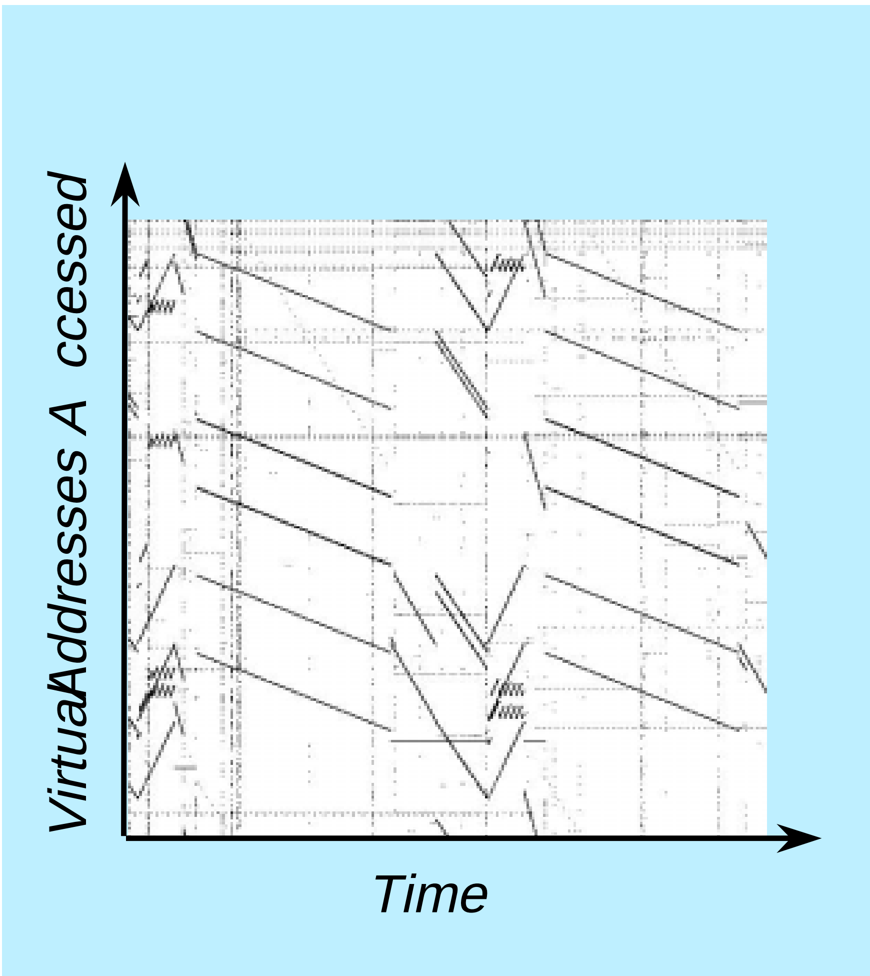

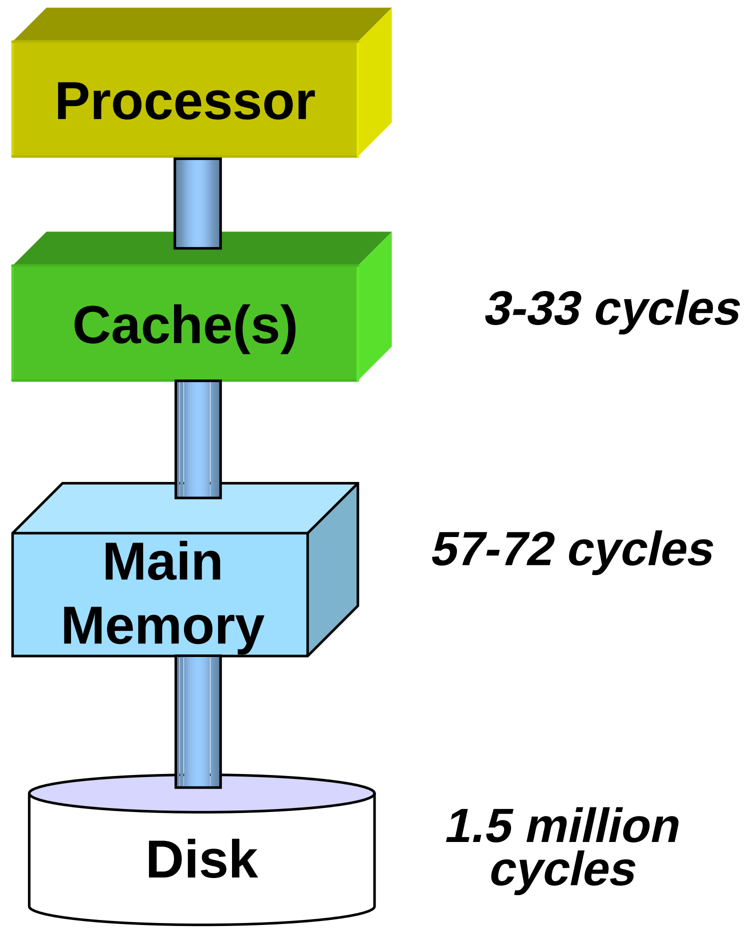

Data-Starved Defense Applications

Many Defense Applications Have Large Data Sets That Are Accessed Non-Contiguously

Result: *These data intensive defense applications perform at 1% of peak.*

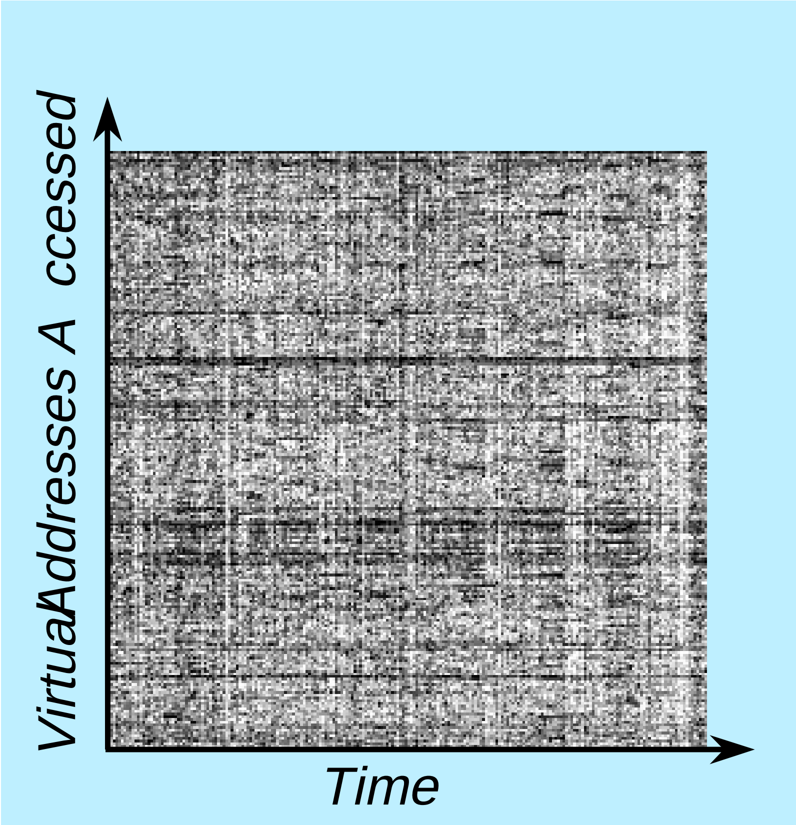

Data Starved Applications

- Radar Cross-Section Modeling (FMM)
- High-Definition Imaging (HDI)
- Terrain Masking
- Relational Databases
- Object-Oriented Databases(OODB)
- Structural Dynamics
- Circuit Simulation

Data Intensive Systems Goal

Enable new regime of realtime data-intensive applications
Hours → Seconds

Performance will be demonstrated by computing 64 looks/min. at a T-72 (tank) using model based ATR.

64 looks today at a T-72 would require 18.5 hours!

Approach

Two complementary tasks:

- **In Situ Processing**

  Logic within memory chips manipulates data within the memory subsystem

- **Adaptive Cache Management**

  Applications manage memory hierarchy so data placement and flow is tailored to application specific needs

Data Intensive Approach: In Situ Processing

Put a Modicum of Intelligence DRAM Chips

- Search and sort in situ
- Execute operations at the site of the data

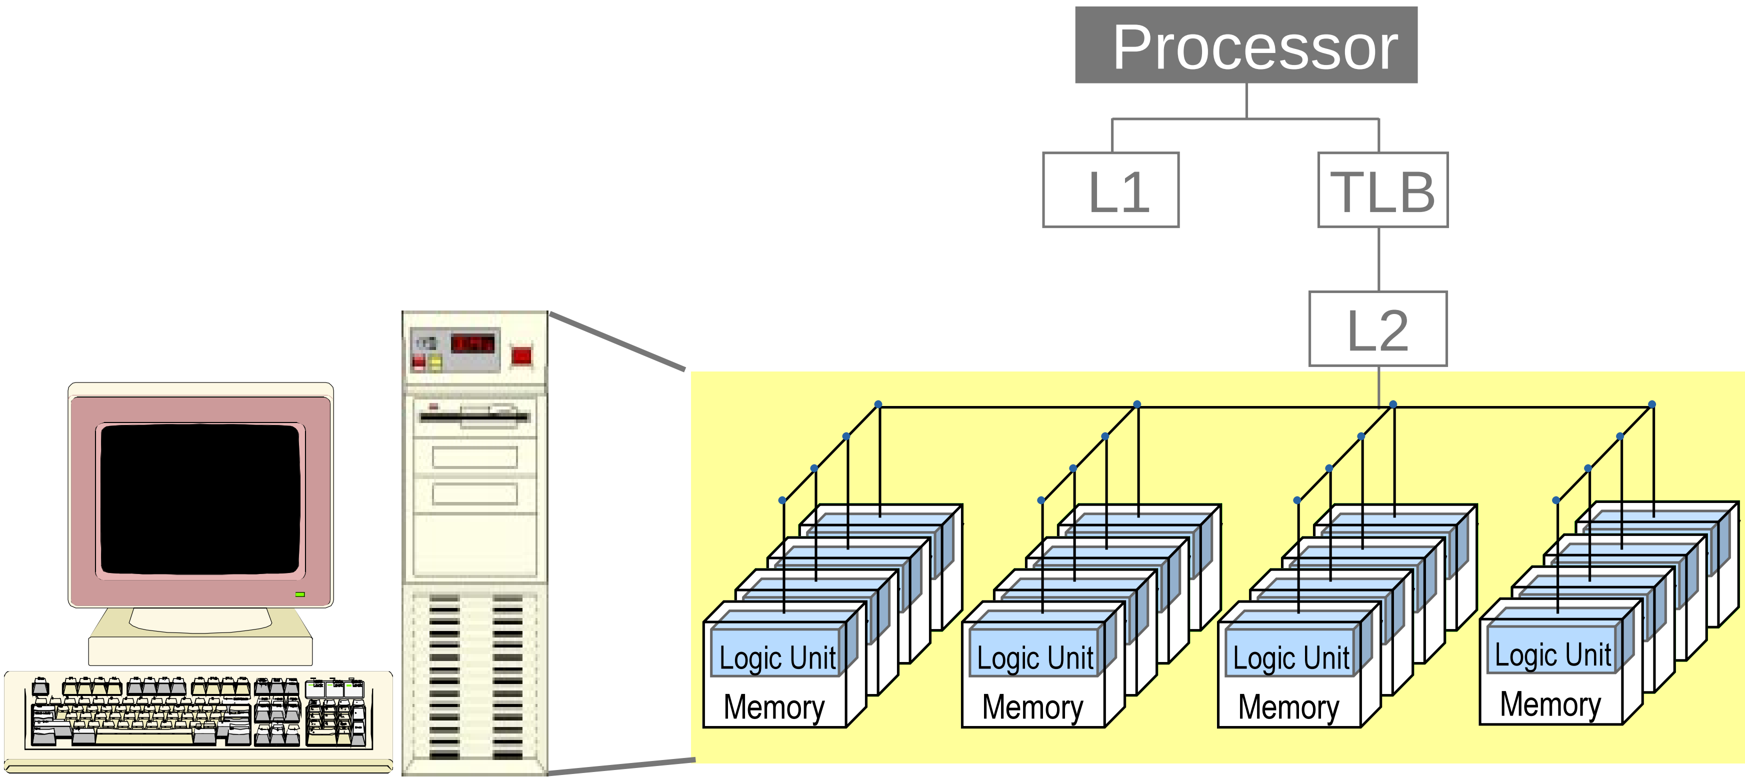

Integration of Logic in Memory

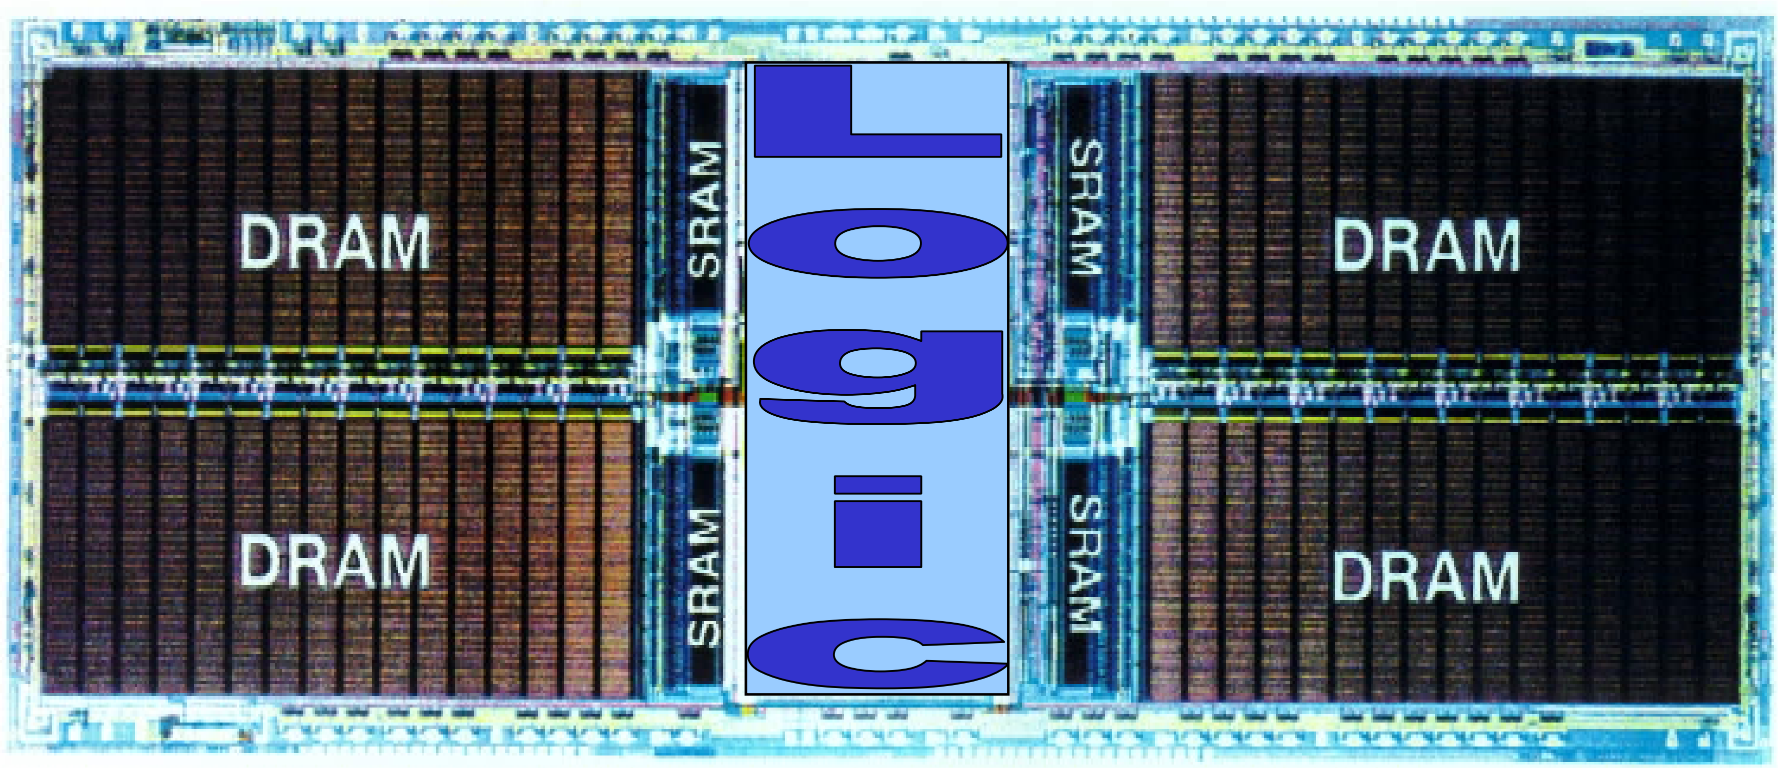

Data Intensive Approach: Adaptive Cache Management

Empower applications with the ability to manage the memory caches

- ***data movement is tailored to their specific needs***

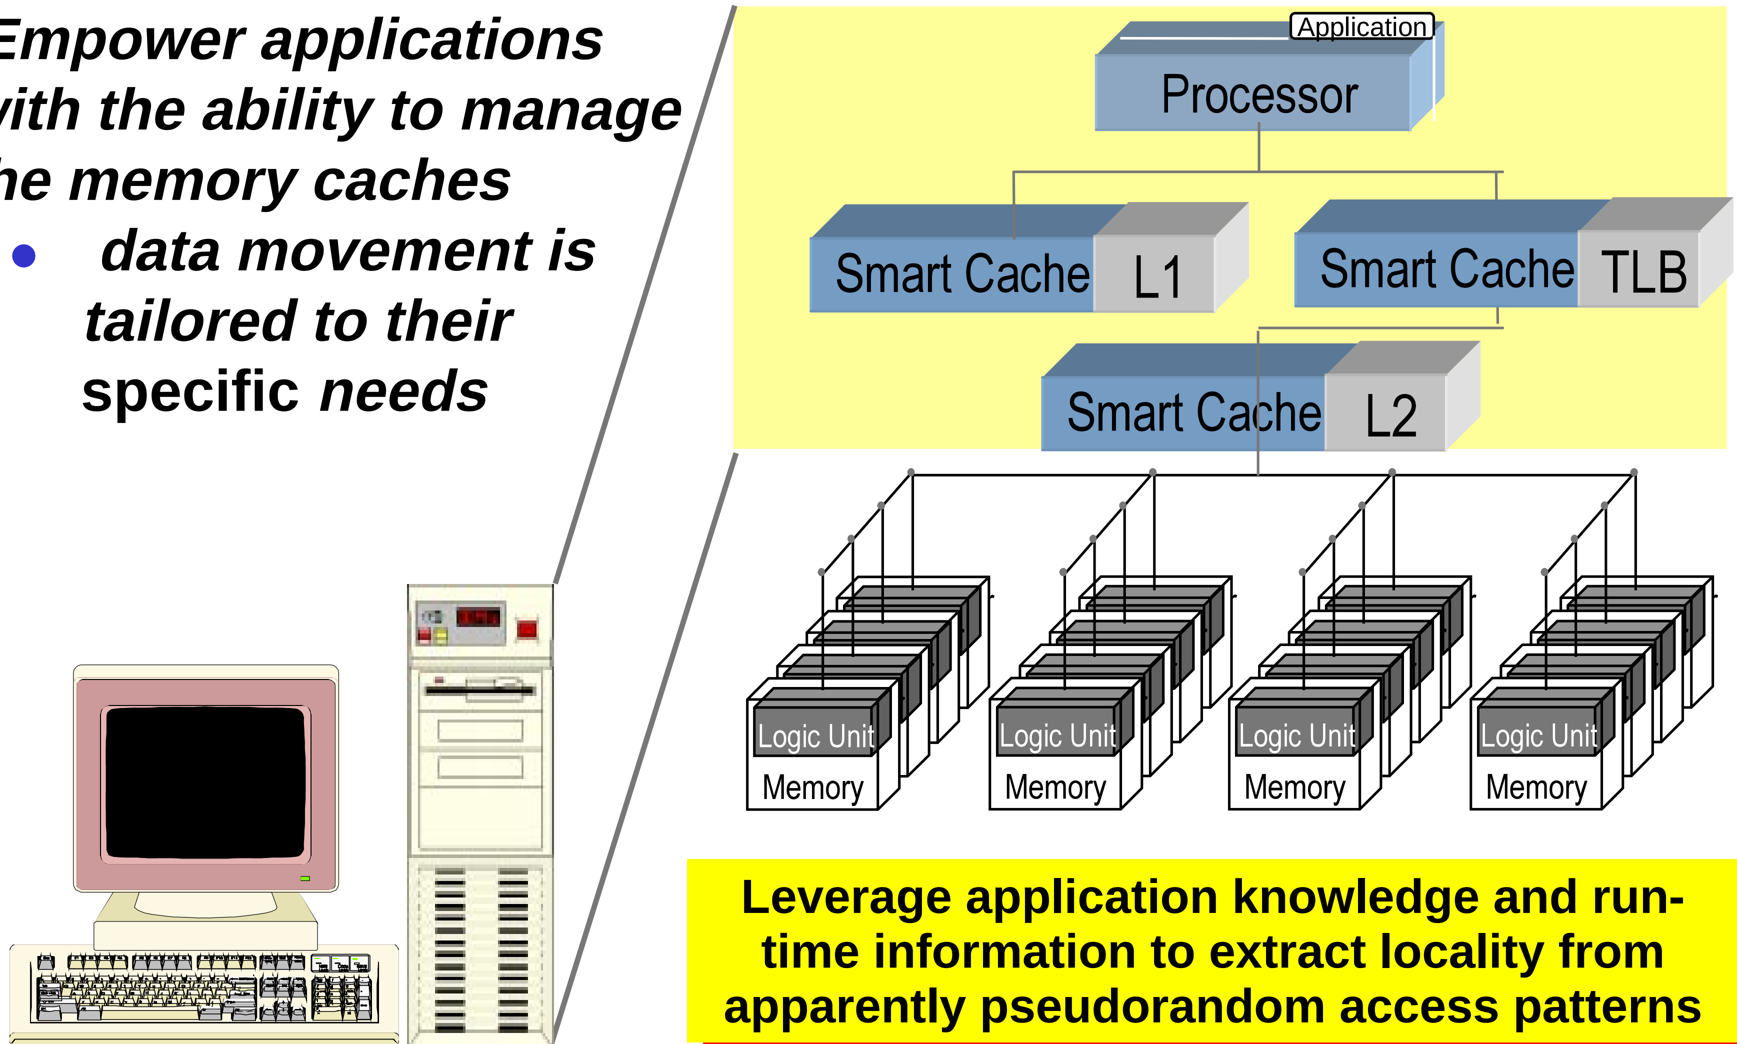

Leverage application knowledge and run-time information to extract locality from apparently pseudorandom access patterns

PRODUCTS: Technology

- **Workstation with new memory hierarchy**
  - **Capable of processing data within memory**
  - **Adaptive caches that allow applications to manage data movement**
  - **Compiler enhancements to leverage above features**
  - **Integration with mainstream OS and database packages**

PRODUCTS: Accelerating Model-Based ATR

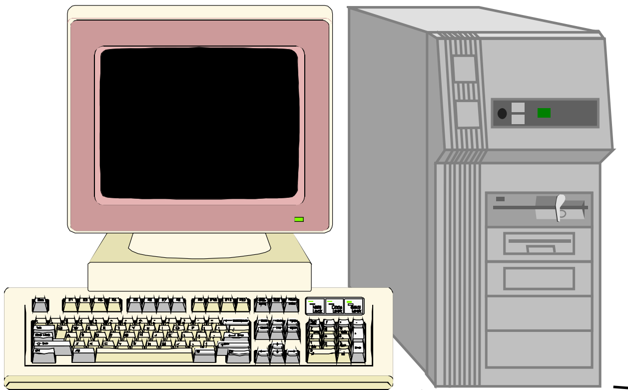

- Compute 64 looks at a T-72 in one minute (1,000,000 ray-patch intersections/second/chip).
- Demonstrate a 20-fold improvement in the performance of the MoM FMM-based radar cross-section modeling code.

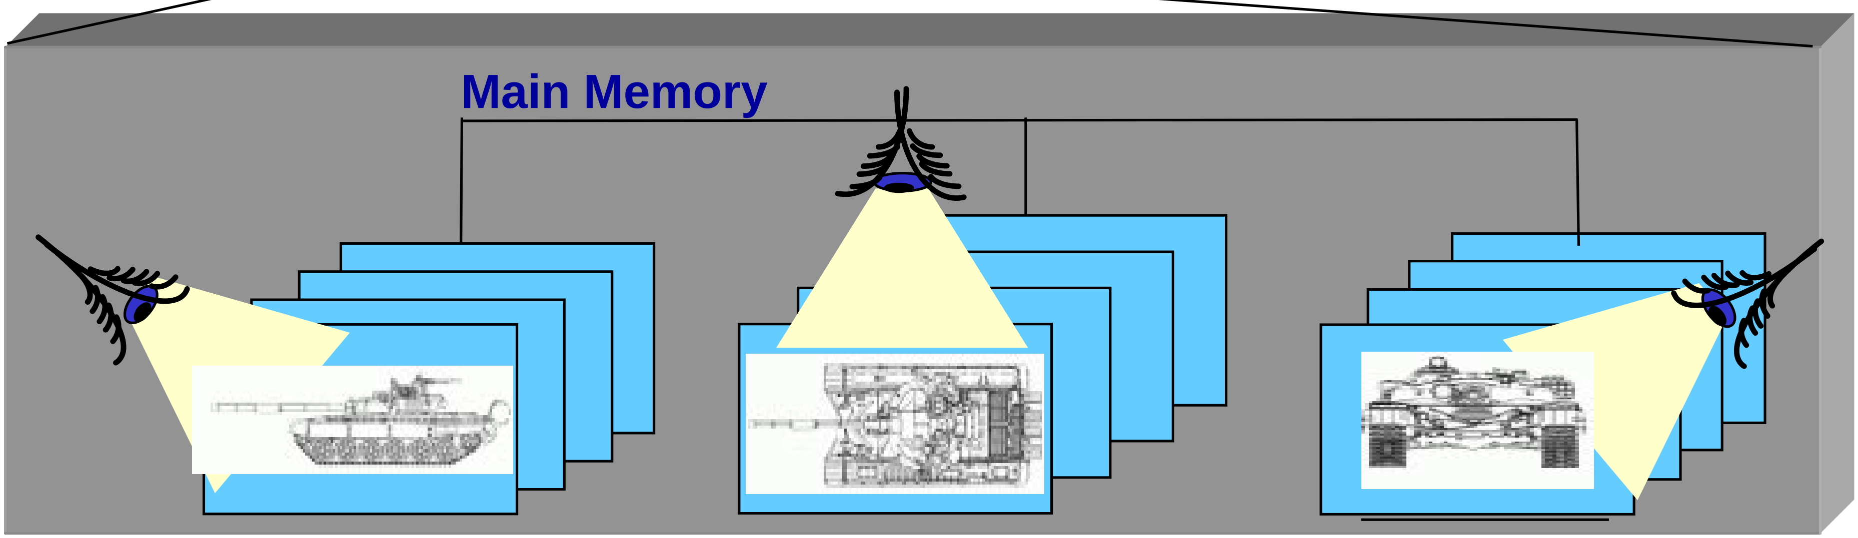

Broad Impact On Many Data Intensive Operations

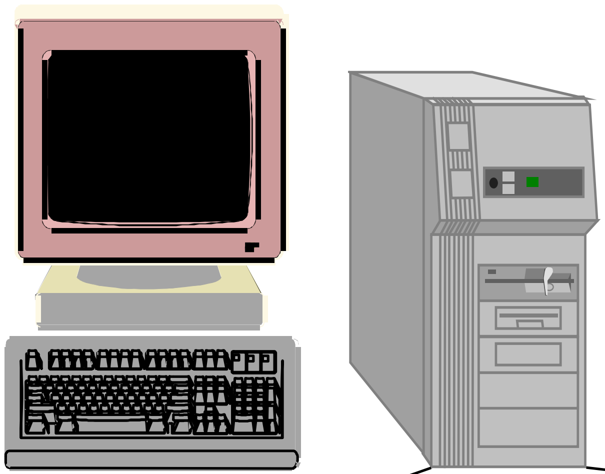

- Search and sort in parallel
- Execute object oriented, methods at the site of the object's data
- Automatic garbage collection, compacting and dereferencing

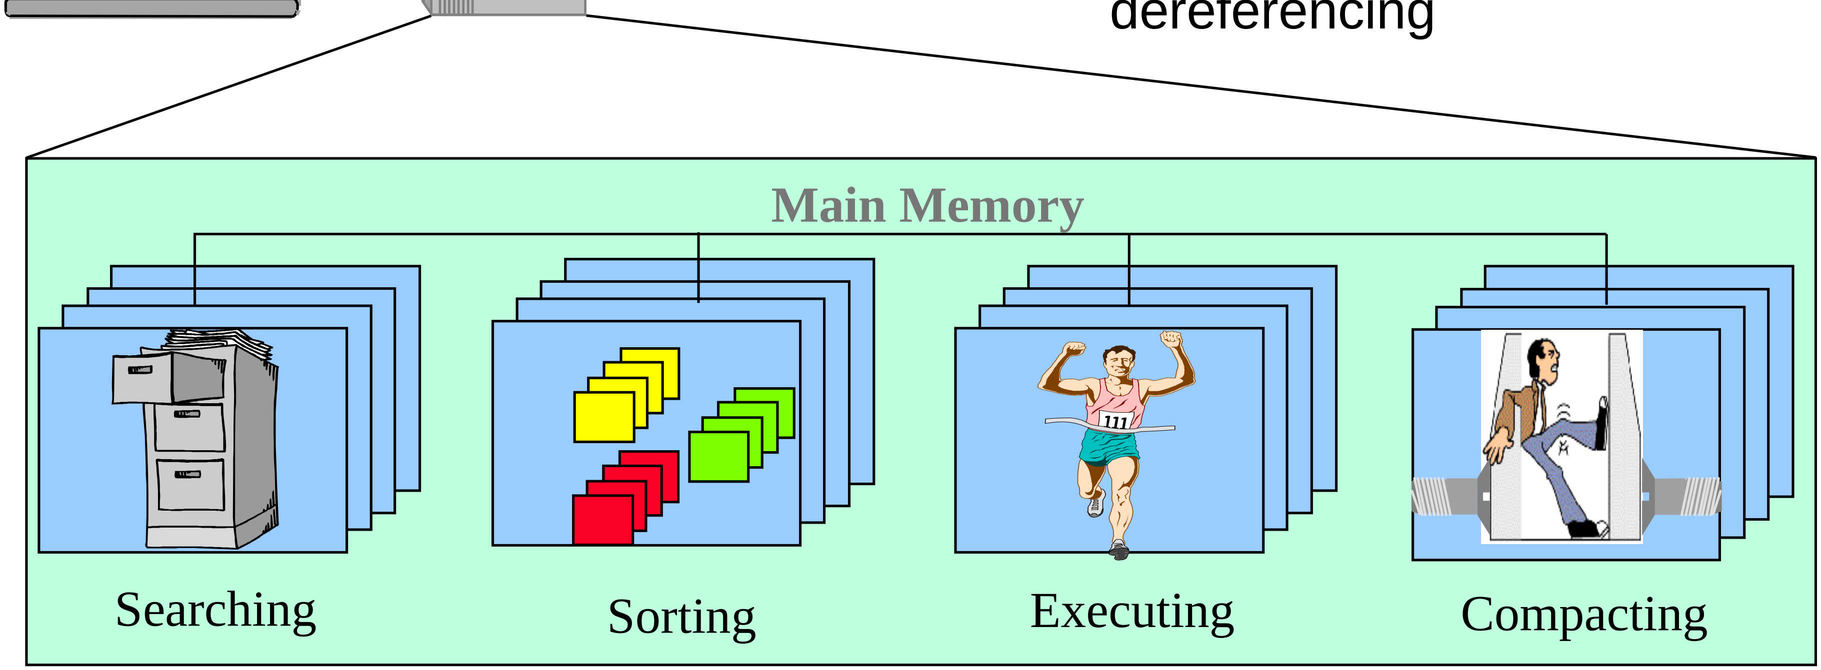

Data Intensive Systems Roadmap